

An Algorithmic Analysis of the Honey-Bee Game*

RUDOLF FLEISCHER[†]

GERHARD J. WOEGINGER[‡]

Abstract

The HONEY-BEE game is a two-player board game that is played on a connected hexagonal colored grid or (in a generalized setting) on a connected graph with colored nodes. In a single move, a player calls a color and thereby conquers all the nodes of that color that are adjacent to his own current territory. Both players want to conquer the majority of the nodes. We show that winning the game is PSPACE-hard in general, NP-hard on series-parallel graphs, but easy on outerplanar graphs.

In the solitaire version, the goal of the single player is to conquer the entire graph with the minimum number of moves. The solitaire version is NP-hard on trees and split graphs, but can be solved in polynomial time on co-comparability graphs.

Keywords: combinatorial game; computational complexity; graph problem.

1 Introduction

The HONEY-BEE game is a popular two-player board game that shows up in many different variants and at many different places on the web (the game is best played on a computer). For a playable version we refer the reader for instance to Axel Born's web-page [1]; see Fig. 1 for a screenshot. The playing field in HONEY-BEE is a grid of hexagonal honey-comb cells that come in various colors; the coloring changes from game to game. The playing field may be arbitrarily shaped and may contain holes, but must always be connected. In the beginning of the game, each player controls a single cell in some corner

*RF acknowledges support by the National Natural Science Foundation of China (No. 60973026), the Shanghai Leading Academic Discipline Project (project number B114), the Shanghai Committee of Science and Technology of China (09DZ2272800), and the Robert Bosch Foundation (Science Bridge China 32.5.8003.0040.0). GJW acknowledges support by the Netherlands Organisation for Scientific Research (NWO grant 639.033.403), and by BSIK grant 03018 (BRICKS: Basic Research in Informatics for Creating the Knowledge Society).

[†]School of Computer Science, IIPL, Fudan University, Shanghai 200433, China. Email: rudolf@fudan.edu.cn.

[‡]gwoegi@win.tue.nl. Department of Mathematics and Computer Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, Netherlands.




Figure 1: Born’s “Biené”. The human player (starting from the top-left corner) is on the edge of losing against the computer (starting from the bottom-right corner).

of the playing field. Usually, the playing area is symmetric and the two players face each other from symmetrically opposing starting cells. In every move a player may call a color c , and thereby gains control over all connected regions of color c that have a common border with the area already under his control. The only restriction on c is that it cannot be one of the two colors used by the two players in their last move before the current move, respectively. A player wins when he controls the majority of all cells. On Born’s web-page [1] one can play against a computer, choosing from four different layouts for the playing field. The computer uses a simple greedy strategy: “Always call the color c that maximizes the immediate gain.” This strategy is short-sighted and not very strong, and an alert human player usually beats the computer after a few practice matches.

In this paper we perform a complexity study of the HONEY-BEE game when played by two players on some arbitrary connected graph instead of the hex-grid of the original game. We will show in Section 4 that HONEY-BEE-2-PLAYERS is NP-hard even on series-parallel graphs, and that it is PSPACE-complete in general. On outerplanar graphs, however, it is quite easy to compute a winning strategy.

In the *solitaire* (single-player) version of HONEY-BEE the goal is to conquer the entire playing field as quickly as possible. Intuitively, a good strategy for the solitaire game




Figure 2: Summary of the complexity results for HONEY-BEE-SOLITAIRE. NP-complete problems have a solid frame, polynomial-time solvable problems have a dashed frame. The results for the graph classes in the three colored boxes imply all other results.

will be close to a strong heuristic for the two-player game. For the solitaire version, our results draw a sharp separation line between easy and difficult cases. In particular, we show in Section 3 that HONEY-BEE-SOLITAIRE is NP-hard for split graphs and for trees, but polynomial-time solvable on co-comparability graphs (which include interval graphs and permutation graphs). Thus, the complexity of the game is well-characterized for the class and subclasses of perfect graphs; see Fig. 2 for a summary of our results.

2 Definitions

We model HONEY-BEE in the following graph-theoretic setting. The playing field is a connected, simple, loopless, undirected graph $G = (V, E)$. There is a set C of k colors, and every node $v \in V$ is colored by some color $col(v) \in C$; we stress that this coloring does not need to be proper, that is, there may be edges $[u, v] \in E$ with $col(u) = col(v)$. For a color $c \in C$, the subset $V_c \subseteq V$ contains the nodes of color c . For a node $v \in V$ and a color $c \in C$, we define the *color- c -neighborhood* $\Gamma(v, c)$ as the set of nodes in V_c either adjacent to v or connected to v by a path of nodes of color c . Similarly, we denote by $\Gamma(W, c) = \bigcup_{w \in W} \Gamma(w, c)$ the color- c -neighborhood of a subset $W \subseteq V$. For a subset $W \subseteq V$ and a sequence $\gamma = \langle \gamma_1, \dots, \gamma_b \rangle$ of colors in C , we define a corresponding sequence of node sets $W_1 = W$ and $W_{i+1} = W_i \cup \Gamma(W_i, \gamma_i)$, for $1 \leq i \leq b$. We say that sequence γ started on W *conquers* the final node set W_{b+1} in b moves, and we denote this situation by $W \rightarrow_\gamma W_{b+1}$. The nodes in $V - W_{b+1}$ are called *free* nodes.

In the *solitaire* version of HONEY-BEE, the goal is to conquer the entire playing field with the smallest possible number of moves. Note that HONEY-BEE-SOLITAIRE is trivial in the case of only two colors. But as we will see in Section 3, the case of three colors can already be difficult.

Problem HONEY-BEE-SOLITAIRE

Input: A graph $G = (V, E)$; a set C of k colors and a coloring $col : V \rightarrow C$; a start node $v_0 \in V$; and a bound b .

Question: Does there exist a color sequence $\gamma = \langle \gamma_1, \dots, \gamma_b \rangle$ of length b such that $\{v_0\} \rightarrow_\gamma V$?

In the *two-player* version of HONEY-BEE, the two players A and B start from two distinct nodes a_0 and b_0 and then extend their regions step by step by alternately calling colors. Player A makes the first move. One round of the game consists of a move of A followed by a move of B . Consider a round, where at the beginning the two players control node sets W_A and W_B , respectively. If player A calls color c , then he extends his region W_A to $W'_A = W_A \cup (\Gamma(W_A, c) - W_B)$. If afterwards player B calls color d , then he extends his region W_B to $W'_B = W_B \cup (\Gamma(W_B, d) - W'_A)$. Note that once a player controls a node, he can never lose it again.

The game terminates as soon as one player controls more than half of all nodes. This player wins the game. To avoid draws, we require that the number of nodes is odd. There are three important rules that constrain the colors that a player is allowed to call.

- R1. A player must never call the color that has just been called by the other player.
- R2. A player must never call the color that he has called in his previous move.
- R3. A player must always call a color that strictly enlarges his territory, unless rules R1 and R2 prevent him from doing so.

What is the motivation for these three rules? Rule R1 is a technical condition that arises from the graphical implementation [1] of the game: Whenever a player calls a color c , his current territory is entirely recolored to color c . This makes it visually easier to recognize the territories controlled by both players. Rule R2 prevents the players from permanently blocking some color for the opponent. Fig. 3 shows a situation where rule R2 actually prevents the game from stalling. Rule R3 is quite delicate, and is justified by situations as depicted in Fig. 4. Rule R3 guarantees that every game must terminate with either a win for player A or a win for player B . Note that rule R2 is redundant except in the case when a player has no move to gain territory (see Fig. 3).




Figure 3: Player A (circled nodes) is leading with four captured nodes over player B (squared nodes) with only two captured nodes. Player B would next like to play black to capture all the white nodes in the next move. Without rule R2, player A could prevent this by repeatedly playing black.




Figure 4: Player A who controls the black node at the left end of the path loses if he calls dark-gray (and hence prefers to call white, light-gray, and black). Player B who controls the white node at the other end of the path loses if he calls light-gray (and hence prefers to call colors white, dark-gray, and black). Rule R3 forces the players to move into the unoccupied territory.

Problem HONEY-BEE-2-PLAYERS

Input: A graph $G = (V, E)$ with an odd number of nodes; a set C of colors and a coloring $col : V \rightarrow C$; two start nodes $a_0, b_0 \in V$.

Question: Can player A enforce a win when the game is played according to the above rules?

Note that HONEY-BEE-2-PLAYERS is trivial in the case of only three colors: The players do not have the slightest freedom in choosing their next color, and always must call the unique color allowed by rules R1 and R2. However we will see in Section 4 that the case of four colors can already be difficult.

Finally we observe that calling a color c always conquers all connected components induced by V_c that are adjacent to the current territory. Hence an equivalent definition of the game could use a graph with node weights (that specify the size of the corresponding connected component) and a *proper* coloring of the nodes. Any instance under the original definition can be transformed into an equivalent instance under the new definition by contracting each connected component of V_c , for some c , into a single node of weight $|V_c|$. However, we are interested in restrictions of the game to particular graph classes, some of which are not closed under edge contractions (as for instance the hex-grid graph of the original HONEY-BEE game).

3 The Solitaire Game

In this section we study the complexity of finding optimally short color sequences for HONEY-BEE-SOLITAIRE. We will show that this is easy for co-comparability graphs, while it is NP-hard for trees and split graphs. Since the family of co-comparability graphs contains interval graphs, permutation graphs, and co-graphs as sub-families, our positive result for co-comparability graphs implies all other positive results in Fig. 2.

A first straightforward observation is that HONEY-BEE-SOLITAIRE lies in NP: Any connected graph $G = (V, E)$ can be conquered in at most $|V|$ moves, and hence such a sequence of polynomially many moves can serve as an NP-certificate.

3.1 The Solitaire Game on Co-Comparability Graphs

A *co-comparability graph* $G = (V, E)$ is an undirected graph whose nodes V correspond to the elements of some partial order $<$ and whose edges E connect any two elements that are incomparable in that partial order, i.e., $[u, v] \in E$ if neither $u < v$ nor $v < u$ holds. For simplicity, we identify the nodes with the elements of the partial order. Golumbic et al. [3] showed that co-comparability graphs are exactly the intersection graphs of continuous real-valued functions over some interval I . If two function curves intersect, the corresponding elements are incomparable in the partial order; otherwise, the curve that lies complete above the other one corresponds to the larger element in the partial order. The function graph representation readily implies that the class of co-comparability graphs is closed

under edge contractions. Therefore, we may w.l.o.g. restrict our analysis of HONEY-BEE-SOLITAIRE to co-comparability graphs with a proper node coloring, i.e., adjacent nodes have distinct colors (in the solitaire game we do not care about the weight of a node after an edge contraction). In this case, every color class is totally ordered because incomparable node pairs have been contracted.

Consider an instance of HONEY-BEE-SOLITAIRE with a minimal start node v_0 (in the partial order on V); a maximal start node could be handled similarly. The function graph representation implies the following observation.

Observation 3.1 *Conquering a node will simultaneously conquer all smaller nodes of the same color.* \square

For any color c , let $Max(c)$ denote the largest node of color c . By Obs. 3.1, it suffices to find the shortest color sequence conquering all nodes $Max(c)$, for all colors c . We can do that by a simple shortest path computation. We assign every node $Max(c)$ weight 0, and all other nodes weight 1. Then we compute a shortest path (with respect to the node-weights) from v_0 to every node $Max(c)$ that is a *maximal element* in the partial order (which is actually exactly the set of all maximal elements). Let OPT denote the smallest cost over all such paths.

For a color sequence $\gamma = \langle \gamma_1, \dots, \gamma_b \rangle$, we define the *length* of γ as $|\gamma| = b$. We also define the *essential length* $ess(\gamma)$ of γ as $|\gamma|$ minus the number of steps where γ conquers a maximal node $Max(c)$ of some color class c . Obviously, $|\gamma| = ess(\gamma) + k$. Note that OPT is the minimal essential cost of any color sequence conquering one of the maximal nodes.

Lemma 3.2 *The optimal solution for HONEY-BEE-SOLITAIRE has cost $OPT + k$.*

Proof. Let γ be a shortest color sequence conquering the entire graph starting at v_0 . After conquering v , γ only needs to conquer all free nodes $Max(c)$ to conquer the entire graph. Thus, $|\gamma| = ess(\gamma) + k \geq OPT + k$. \square

Theorem 3.3 *HONEY-BEE-SOLITAIRE starting at an extremal node v_0 can be solved in polynomial time on co-comparability graphs.*

Proof. Given the co-comparability graph G , we can compute the underlying partial order $<$ in polynomial time [3]. Assigning the weights and solving one single source shortest path problem starting at v_0 also takes polynomial time. \square

We can also formulate this algorithm as a dynamic program. For any node v , let $D(v)$ denote the essential length of the shortest color sequence γ that can conquer v when

starting at v_0 . For any color c , let $\min_v(c)$ denote the smallest node of color c connected to v , if such nodes exist. Then we can compute $D(v)$ recursively as follows:

$$D(v_0) = 0$$

and

$$D(v) = \min_c (D(\min_v(c)) + \delta_v) ,$$

where $D(\min_v(c)) = \infty$ if $\min_v(c)$ is undefined, and $\delta_v = 0$ (1) if v is (not) a maximal node for some color class.

Clearly, this dynamic program simulates the shortest path computation of our first algorithm and we have $OPT = \min_v (D(v) + k)$, where we minimize over all maximal nodes v . We now extend the dynamic program to the case that v_0 is not an extremal element. The problem is that we now must extend our territory in two directions. If we choose a move that makes good progress upwards it may make little progress downwards, or vice versa. In particular, the optimal strategy cannot be decomposed into two independent optimal strategies, one conquering upwards and one conquering downwards. Analogously to the algorithm above, for a color c define $\min_v(c)$ as the smallest node of color c , and $\max_v(c)$ as the largest node of color c connected to a node v .

Unfortunately, we must now redefine the essential length of a color sequence γ . In our original definition, we did not count coloring steps that conquered maximal elements of some color class. This is intuitively justified by the fact that these steps must be done by any color sequence conquering the entire graph at some time, therefore it is advantageous to do them as early as possible (which is guaranteed by giving these moves cost 0). But now we must also consider the minimal nodes of each color class. An optimal sequence conquering the entire graph will at some time have conquered a minimal node and a maximal node. Afterwards, it will only call extremal nodes for some color class. If both extremal nodes of a color class are still free, we only need *one* move to conquer both simultaneously. If one of them had been captured earlier, we still need to conquer the other one. This indicates that we should charge 1 for the first extremal node conquered while the second one should be charged 0, as before. If both nodes are conquered in the same move, we should also charge 0. Therefore, we now define the *essential length* $ess(\gamma)$ of γ as $|\gamma|$ minus the number of steps where γ conquers the second extremal node of some color class.

For a node v below v_0 or incomparable to v_0 and a node w above v_0 or incomparable to v_0 let $D(v, w)$ denote the essential length of the shortest color sequence γ that can conquer v and w when starting at v_0 . Note that we do not need to keep track of which first extremal nodes of a color class have been conquered because we can deduce this from the two nodes v and w currently under consideration. In particular, we can compute $D(v, w)$ recursively as follows:

$$D(v_0, v_0) = 0$$

and

$$D(v, w) = \min_c (D(v, \min_w(c)) + \delta_w(v), D(\max_v(c), w) + \delta_v(w)) ,$$

where $\delta_v(w) = 0$ if and only if w is an extremal node of some color class c and the other extremal node of color class c is either between v and w , or incomparable to either v or w , or both (it was either conquered earlier, or it will be conquered in this step); otherwise, $\delta_v(w) = 1$. Obviously, $|\gamma| = \text{ess}(\gamma) + k$.

Lemma 3.4 *The optimal solution for HONEY-BEE-SOLITAIRE has cost $\min_{v,w} (D(v,w) + k)$, where we minimize over all minimal nodes v and all maximal nodes w .*

Proof. Let γ be a shortest color sequence conquering the entire graph starting at v_0 . Let v be the first minimal node conquered by γ and w the first maximal node. After conquering v and w , γ only needs to conquer all free nodes $\text{Max}(c)$ to conquer the entire graph. Thus, $|\gamma| \geq D(v,w) + k$. \square

Theorem 3.5 *HONEY-BEE-SOLITAIRE can be solved in polynomial time on co-comparability graphs.* \square

3.2 The Solitaire Game on Split Graphs

A *split graph* is a graph whose node set can be partitioned into an induced clique and into an induced independent set. We will show that HONEY-BEE-SOLITAIRE is NP-hard on split graphs. Our reduction is from the NP-hard **Feedback Vertex Set** (FVS) problem in directed graphs; see for instance Garey and Johnson [2].

Problem FVS

Input: A directed graph (X, A) ; a bound $t < |X|$.

Question: Does there exist a subset $X' \subseteq X$ with $|X'| = t$ such that the directed graph induced by $X - X'$ is acyclic?

Theorem 3.6 *HONEY-BEE-SOLITAIRE on split graphs is NP-hard.*

Proof. Consider an instance (X, A, t) of FVS. To construct an instance (V, E, b) of HONEY-BEE-SOLITAIRE, we first build a clique from the nodes in X together with a new node v_0 , the start node of HONEY-BEE-SOLITAIRE, where each node $x \in X + v_0$ has a different color c_x . Next, we build the independent set. For every arc $(x, y) \in A$, we introduce a corresponding node $v(x, y)$ of color c_y which is only connected to node x in the clique, i.e., it has degree one. Finally, we set $b = |X| + t$. We claim that the constructed instance of HONEY-BEE-SOLITAIRE has answer YES, if and only if the instance of FVS has answer YES.

Assume that the FVS instance has answer YES. Let X' be a smallest feedback set whose removal makes (X, A) acyclic. Let π be a topological order of the nodes in $X - X'$,

and let τ be an arbitrary ordering of the nodes in X' . Consider the color sequence γ of length $|X| + t$ that starts with τ , followed by π , and followed by τ again. We claim that $\{v_0\} \rightarrow_\gamma V$. Indeed, γ first runs through τ and π and thereby conquers all clique nodes. Every independent set node $v(x, y)$ with $y \in X'$ is conquered during the second transversal of τ . Every independent set node $v(x, y)$ with $y \in X - X'$ is conquered during the transversal of π , since π first conquers x with color c_x , and afterwards $v(x, y)$ with color y .

Next assume that the instance of HONEY-BEE-SOLITAIRE has answer YES. Let γ be a color sequence of length $b = |X| + t$ conquering V . Define X' as the set of nodes x such that color c_x occurs at least twice in γ ; clearly, $|X'| \leq t$. Consider an arc $(x, y) \in A$ with $x, y \in X - X'$. Since γ contains color c_y only once, it must conquer node $v(x, y)$ of color c_y after node $v(x)$ of color c_x . Hence, γ induces a topological order of $X - X'$. \square

The construction in the proof above uses linearly many colors. What about the case of few colors? On split graphs, HONEY-BEE-SOLITAIRE can always be solved by traversing the color set C twice; the first traversal conquers all clique nodes, and the second traversal conquers all remaining free independent set nodes. Thus, every split graph can be completely conquered in at most $2|C|$ steps. If there are only few colors, we can simply check all color sequences of this length $2|C|$.

Theorem 3.7 *If the number of colors is bounded by a fixed constant, HONEY-BEE-SOLITAIRE on split graphs is polynomial-time solvable.* \square

3.3 The Solitaire Game on Trees

In this section we will show that HONEY-BEE-SOLITAIRE is NP-hard on trees, even if there are at only three colors. We reduce HONEY-BEE-SOLITAIRE from a variant of the **Shortest Common Supersequence** (SCS) problem which is known to be NP-complete (see Middendorf [4]).

Problem SCS

Input: A positive integer t ; finite sequences $\sigma_1, \dots, \sigma_s$ with elements from $\{0, 1\}$ with the following properties: (i) All sequences have the same length. (ii) Every sequence contains exactly two 1s, and these two 1s are separated by at least one 0.

Question: Does there exist a sequence σ of length t that contains $\sigma_1, \dots, \sigma_s$ as subsequences?

Middendorf's hardness result also implies the hardness of the following variant of SCS:

Problem Modified SCS (MSCS)

Input: A positive integer t ; finite sequences $\sigma_1, \dots, \sigma_s$ with elements from $\{0, 1, 2\}$ with the following property: In every sequence any two consecutive elements are distinct, and no sequence starts with 2.

Question: Does there exist a sequence σ of length t that contains $\sigma_1, \dots, \sigma_s$ as subsequences?

Theorem 3.8 *MSCS is NP-complete.*

Proof. Here is a reduction from SCS to MSCS. Consider an arbitrary sequence τ with elements from $\{0, 1\}$. We define $f(\tau)$ as the sequence we obtain from replacing every occurrence of the element 0 in τ by two consecutive elements 0 and 2. Now consider an instance $(\sigma_1, \dots, \sigma_s, t)$ of SCS. We construct an instance $(\sigma'_1, \dots, \sigma'_s, t')$ of MSCS by setting $\sigma'_i = f(\sigma_i)$, for $1 \leq i \leq s$. Then, for any sequence σ with elements from $\{0, 1\}$, σ is a common supersequence of $\sigma_1, \dots, \sigma_s$ if and only if $f(\sigma)$ is a common supersequence of $\sigma'_1, \dots, \sigma'_s$. This implies the NP-hardness of MSCS. \square

Theorem 3.9 *HONEY-BEE-SOLITAIRE is NP-hard on trees, even in case of only three colors.*

Proof. We reduce MSCS to HONEY-BEE-SOLITAIRE on trees. Consider an instance $(\sigma_1, \dots, \sigma_s, t)$ of MSCS. We use color set $C = \{0, 1, 2\}$. We first construct a root v_0 of color 2. Then we attach a path of length $|\sigma_i|$ to v_0 for each sequence σ_i , where an element j is colored j . See the left half of Fig. 6 for an example. Finally, we set $b = t$. It is straightforward to see that the constructed instance of HONEY-BEE-SOLITAIRE has answer YES if and only if the instance of MSCS has answer YES. \square

4 The Two-Player Game

In this section we study the complexity of the two-player game. While on outerplanar graphs the players can compute their winning strategies in polynomial time, this problem is NP-hard for series-parallel graphs with four colors, and PSPACE-complete with four colors on arbitrary graphs. Our positive result for outerplanar graphs works for an arbitrary number of colors. Our negative results work for four colors, which is the strongest possible type of result (recall that instances with three colors are trivial to solve).

4.1 The Two-Player Game on Outer-Planar Graphs

A graph is *outer-planar* if it contains neither K_4 nor $K_{2,3}$ as a minor. Outer-planar graphs have a planar embedding in which every node lies on the boundary of the so-called *outer face*. For example, every tree is an outer-planar graph.




Figure 5: An outerplanar graph with start nodes a_0 and b_0 . Player A (circled nodes) has conquered the light-gray colored nodes, i.e., $U = 2$ and $L = 2$. Eventually, A will also conquer l_1 , since Player B cannot reach it.

Consider an outer-planar graph $G = (V, E)$ as an instance of HONEY-BEE-2-PLAYERS with starting nodes a_0 and b_0 in V , respectively. The starting nodes divide the nodes on the boundary of the outer face F into an upper chain u_1, \dots, u_s and a lower chain ℓ_1, \dots, ℓ_t , where u_1 and ℓ_1 are the two neighbors of a_0 on F , while u_s and ℓ_t are the two neighbors of b_0 on F . We stress that this upper and lower chain are not necessarily disjoint (for instance, articulation nodes will occur in both chains).

Now consider an arbitrary situation in the middle of the game. Let U (respectively L) denote the largest index k such that player A has conquered node u_k (respectively node ℓ_k). See Fig. 5 to illustrate these definitions and the following lemma.

Lemma 4.1 *Let X denote the set of nodes among u_1, \dots, u_U and ℓ_1, \dots, ℓ_L that currently do neither belong to A nor to B . Then no node in X can have a neighbor among $u_{U+1}, \dots, u_s, b_0, \ell_t, \dots, \ell_{L+1}$.*

Proof. The existence of such a node in X would lead to a K_4 -minor in the outer-planar graph. \square

Theorem 4.2 *HONEY-BEE-2-PLAYERS on outer-planar graphs is polynomial-time solvable.*

Proof. The two indices U and L encode all necessary information on the future behavior of player A . Eventually, he will own all nodes u_1, \dots, u_U and ℓ_1, \dots, ℓ_L , and the possible future expansions of his area beyond u_U and ℓ_L only depend on U and L . Symmetric observations hold true for player B .

As every game situation can be concisely described by just four indices, there is only a polynomial number $O(|V|^4)$ of relevant game situations. The rest is routine work in

combinatorial game theory: We first determine the winner for every end-situation, and then by working backwards in time we can determine the winners for the remaining game situations. \square

4.2 The Two-Player Game on Series-Parallel Graphs

A graph is *series-parallel* if it does not contain K_4 as a minor. Equivalently, a series-parallel graph can be constructed from a single edge by repeatedly doubling edges, or removing edges, or replacing edges by a path of two edges with a new node in the middle of the path. We stress that we do not know whether the two-player game on series-parallel graphs is contained in the class NP (and we actually see no reason why it should lie in NP); therefore the following theorem only states NP-hardness.

Theorem 4.3 *For four (or more) colors, problem HONEY-BEE-2-PLAYERS on series-parallel graphs is NP-hard.*

Proof. We use the color set $C = \{0, 1, 2, 3\}$. A central feature of our construction is that player B will have no real decision power, but will only follow the moves of player A : If player A starts a round by calling color 0 or 1, then player B must follow by calling the other color in $\{0, 1\}$ (or waste his move). And if player A starts a round by calling color 2 or 3, then player B must call the other color in $\{2, 3\}$ (or waste his move). In the even rounds the players will call the colors in $\{0, 1\}$ and in the odd rounds they will call the colors in $\{2, 3\}$. Both players are competing for a set of honey pots in the middle of the battlefield, and need to get there as quickly as possible. If a player deviates from the even-odd pattern indicated above, he might perhaps waste his move and delay the game by one round (in which neither player comes closer to the honey pots), but this remains without further impact on the outcome of the game.

The proof is by reduction from the supersequence problem SCS with binary sequences; see Section 3.3. Consider an instance $(\sigma_1, \dots, \sigma_s, t)$ of SCS, and let n denote the common length of all sequences σ_i . We first construct two start nodes a_0 and b_0 of colors 2 and 3, respectively. For each sequence σ_i with $1 \leq i \leq s$ we do the following:

- We construct a path P_i that consists of $2n - 1$ nodes and that is attached to a_0 : The n nodes with odd numbers mimic sequence σ_i , while the $n - 1$ nodes with even numbers along the path all receive color 2. The first node of P_i is adjacent to a_0 , and its last node is connected to a so-called honey pot H_i .
- The *honey pot* H_i is a long path consisting of $4st$ nodes of color 3. Intuitively, we may think of a honey pot as a single node of large weight, because conquering one of the nodes will simultaneously conquer the entire path.
- Every honey pot H_i can also be reached from b_0 by another path Q_i that consists of $2t - 1$ nodes. Nodes with odd numbers get color 0, and nodes with even numbers

4.3 The Two-Player Game on Arbitrary Graphs

In this section we will show that problem HONEY-BEE-2-PLAYERS is PSPACE-complete on arbitrary graphs. Our reduction is from the PSPACE-complete **Quantified Boolean Formula (QBF)** problem; see for instance Garey & Johnson [2].

Problem QBF

Input: A quantified Boolean formula with $2n$ variables in conjunctive normal form: $\exists x_1 \forall x_2 \cdots \exists x_{2n-1} \forall x_{2n} \wedge_j C_j$, where the C_j are clauses of the form $\vee_k l_{jk}$, where the l_{jk} are literals.

Question: Is the formula true?

Theorem 4.4 *For four (or more) colors, problem HONEY-BEE-2-PLAYERS on arbitrary graphs is PSPACE-complete.*

Proof. We reduce from QBF. Let $F = \exists x_1 \forall x_2 \cdots \exists x_{2n-1} \forall x_{2n} \wedge_j C_j$ be an instance of QBF. We construct a bee graph $G_F = (V, E)$ with four colors (white, light-gray, dark-gray, and black) such that player A has a winning strategy if and only if F is true. Let a_0 (colored light-gray) and b_0 (colored dark-gray) denote the start nodes of players A and B , respectively.

Each player controls a *pseudo-path*, that is, a path where some nodes may be duplicated as parallel nodes in a diamond-shaped structure; see Fig. 7. A so-called *choice pair* consists of a node on a pseudo-path together with some duplicated node in parallel. The start nodes are at one end of the respective pseudo-paths, and the players can conquer the nodes on their own path without interference from the other player. However, they must do so in a timely manner because either path ends at a humongous *honey pot*, denoted respectively by H_A and H_B . A honey pot is a large clique of identically-colored nodes (we may think of it as a single node of large weight, because conquering one node will simultaneously conquer the entire clique). Both honey pots have the same size but different colors, namely black (H_A) and white (H_B), and they are connected to each other by an edge. Consequently, both players must rush along their pseudo-paths as quickly as possible to reach their honey pot before the opponent can reach it and to prevent the opponent from winning by conquering both honey pots. The last nodes before the honey pots are denoted by a_f and b_f , respectively. They separate the last variable gadgets (described below) from the honey pots.

Fig. 7 shows an overview of the pseudo-paths and one *variable gadget* in detail. A variable gadget is a part of the two pseudo-paths corresponding to a pair of variables $\exists x_{2i-1} \forall x_{2i}$, for some $i \geq 1$. For player A , the gadget starts at node a_{i-1} with a choice pair a_{2i-1}^F and a_{2i-1}^T , colored white and black, respectively. The first node conquered by A will determine the truth value for variable x_{2i-1} . In the same round, player B has a choice on his pseudo-path P_B between nodes b_{2i-1}^F and b_{2i-1}^T . Since these nodes have the

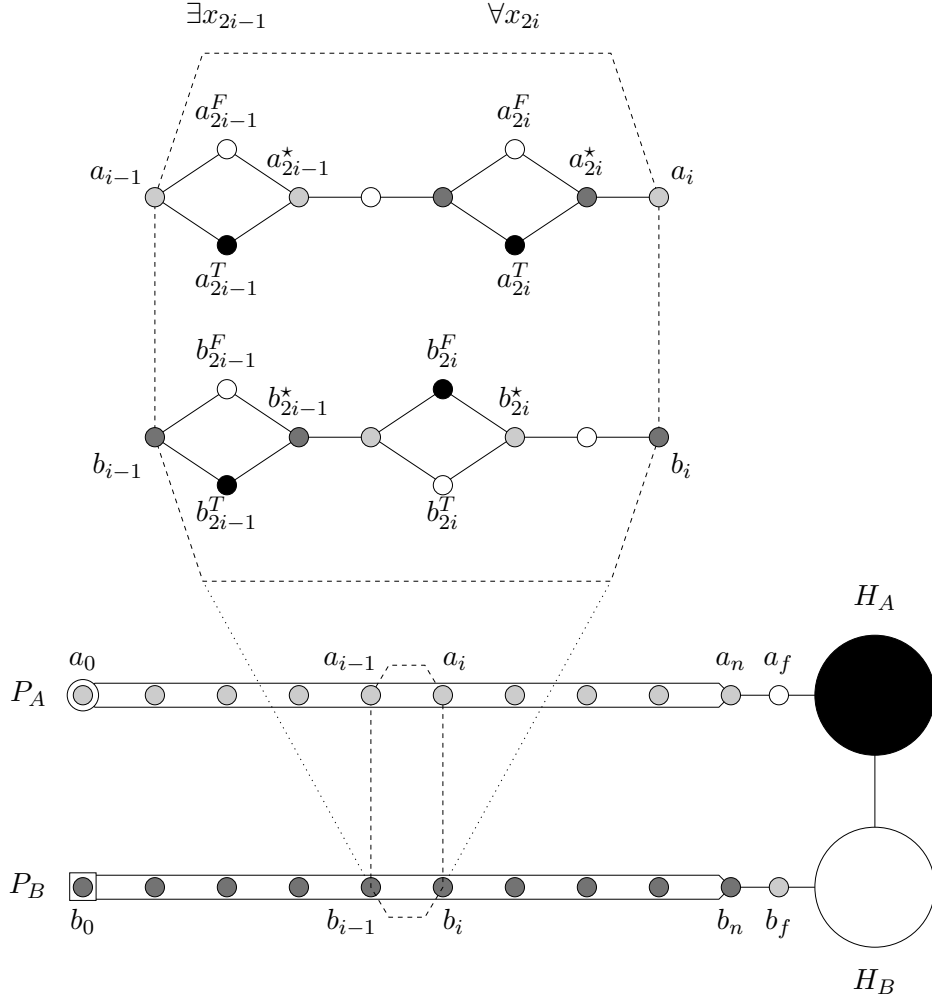


Figure 7: The variable gadget in the proof of Thm. 4.4.

same color as A 's choices in the same round, B actually does not have a choice but must select the other color not chosen by A .

Three rounds later, player B has a choice pair b_{2i}^F and b_{2i}^T , assigning a truth value to variable x_{2i} . In the next step (which is in the next round), player A has a choice pair a_{2i}^F and a_{2i}^T with the same colors as B 's choice pair for x_{2i} . Again, this means that A does not really have a choice but must select the color not chosen by B in the previous step. Since we want A to conquer those clauses containing a literal set to true by player B , the colors in B 's choice pair have been switched, i.e., b_{2i}^F is black and b_{2i}^T is white.

Note that all the nodes a_0, a_1, \dots, a_n are light-gray and all the nodes b_0, b_1, \dots, b_n are dark-gray. This allows us to concatenate as many variable gadgets as needed. Further note that a_f is white, while b_f is light-gray.




Figure 8: The waiting gadgets for existential variables (W_{2i-1}^F and W_{2i-1}^T , the two top paths) and universal variables (W_{2i}^F and W_{2i}^T , the two bottom paths) in the proof of Thm. 4.4. Note that usually one of the two waiting paths W_k^F or W_k^T would be connected to H_j because we may assume that a clause does not contain x_k and $\overline{x_k}$ at the same time.

The *clause* gadgets are very simple. Each clause C_j corresponds to a small honey pot H_j of color white. The size of the small honey pots is smaller than the size of the large honey pots H_A and H_B , but large enough such that player A loses if he misses one of them. Player A should conquer H_j if and only if C_j is true in the assignment chosen by the players while conquering their respective pseudo-paths. We could connect a_{2i-1}^T directly with H_j if C_j contains literal x_{2i-1} , however then player A could in subsequent rounds shortcut his pseudo-path by entering variable gadgets for the other variables in C_j from H_j . To prevent this from happening, we place waiting gadgets between the variable gadgets and the clauses.

Let a_k^* denote the node on P_A right after the choice pair a_k^F and a_k^T , for $k = 1, \dots, 2n$; similarly, b_k^* are the nodes on P_B right after B 's choice pairs. A *waiting gadget* W_k consists of two copies W_k^F and W_k^T of the sub-path of P_A starting at a_k^* and ending at a_n , see Fig. 8. If clause C_j contains literal x_k , H_j is connected to the node w_n^T corresponding to a_n in W_k^T ; if C_j contains literal $\overline{x_k}$, H_j is connected to the node w_n^F corresponding to a_n in W_k^F . If $k = 2i - 1$ (i.e., we have an existential variable x_{2i-1} whose value is assigned by player A), then a_{2i-1}^F and b_{2i-1}^F are connected to w_{2i-1}^{*F} , and a_{2i-1}^T and b_{2i-1}^T are connected to w_{2i-1}^{*T} . If $k = 2i$ (i.e., we have a universal variable x_{2i} whose value is assigned by player B), then a_{2i}^F and b_{2i}^* are connected to w_{2i}^{*F} , and a_{2i}^T and b_{2i-1}^* are connected to w_{2i}^{*T} .

Finally, we connect b_f with all clause honey pots H_j to give player B the opportunity to conquer all those clauses that contain no true literal. This completes the construction of G_F . Fig. 9 shows the complete graph G_F for a small example formula F .

We claim that player A has a winning strategy on G_F if and only if formula F is true. It is easy to verify that player A can indeed win if F is true. All he has to do is to conquer those nodes in his existential choice pairs corresponding to the variable values in a satisfying assignment for F . For the existential variables, he has full control to select any value, and for the universal variables he must pick the opposite color as selected by player B in the previous step, which corresponds to setting the variable to exactly the value that player B has selected. Hence player B can block a move of player A by appropriately selecting a value for a universal variable. Note that no other blocking moves of player B are advantageous: If B blocks A 's next move by choosing a color that does not make progress on his own pseudo-path, then A will simply make an arbitrary waiting move and then in the next round B cannot block A again. When player A conquers node a_n , he will simultaneously conquer the last nodes in all waiting gadgets corresponding to true literals. Since every clause contains a true literal for a satisfying assignment, player A can then in the next round conquer a_f together with all clause honey pots (which all have color white). Player B will respond by conquering b_f , and the game ends with both players conquering their own large honey pots H_A and H_B , respectively. Since player A got all clause honey pots, he wins.

To make this argument work, we must carefully chose the sizes of the honey pots. Each pseudo-path contains $9n + 1$ nodes, of which at most n can be conquered by the other player. The waiting gadgets contain two paths of length $9k + 6$ for existential variables and $9k + 1$ for universal variables. At the end, player A will have conquered one of the two paths completely and maybe some parts of the sibling path, that is, we do not know exactly the final owner of less than n^2 nodes. The clause honey pots should be large enough to absorb this fuzzyness, which means it is sufficient to give them $2n^2$ nodes. The honey pots H_A and H_B should be large enough to punish any foul play by the players, that is, when they do not strictly follow their pseudo-paths. It is sufficient to give them $2n^3$ nodes.

To see that F is true if player A has a winning strategy note that player A must strictly follow his pseudo-path, as otherwise player B could beat him by reaching the large honey pots first. Thus player A 's strategy induces a truth assignment for the existential variables. Similarly, player B 's strategy induces a truth assignment for the universal variables. Player A can only win if he also conquers all clause honey pots, and hence the players must have chosen truth values that make at least one literal per clause true. This means that formula F is satisfiable. \square

5 Conclusions

We have modeled the Honey Bee game as a combinatorial game on colored graphs. For the solitaire version, we have analyzed the complexity on many classes of perfect graphs. For the two player version, we have shown that even the highly restricted case of series-parallel graphs is hard to tackle. Our results draw a clear separating line between easy and hard

variants of these problems.

6 Acknowledgements

Part of this research was done while G. Woeginger visited Fudan University in 2009.

References

- [1] A. Born. Flash application for the computer game “*Biene*” (*Honey-Bee*), 2009.
<http://www.ursulinen.asn-graz.ac.at/Bugs/htm/games/biene.htm>.
- [2] M. R. Garey and D. S. Johnson. *Computers and Intractability — A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.
- [3] M. C. Golumbic, D. Rotem, and J. Urrutia. Comparability graphs and intersection graphs. *Discrete Mathematics*, 43(1):37–46, 1983.
- [4] M. Middendorf. More on the complexity of common superstring and supersequence problems. *Theoretical Computer Science*, 125:205–228, 1994.




Figure 9: The reduction in the proof of Thm. 4.4 would produce this graph for the formula $F = (x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \overline{x_2} \vee x_4) \wedge (\overline{x_2} \vee x_3 \vee \overline{x_4})$.

